When IoT Keeps People in the Loop: A Path Towards a New Global Utility

Vitaly Petrov†, Konstantin Mikhaylov, Dmitri Moltchanov, Sergey Andreev, Gabor Fodor, Johan Torsner, Halim Yanikomeroglu, Markku Juntti, and Yevgeni Koucheryavy

arXiv:1703.00541v2 [cs.NI] 3 Mar 2017

***Abstract*—While the Internet of Things (IoT) has made significant progress along the lines of supporting individual machine-type applications, it is only recently that the importance of people as an integral component of the overall IoT infrastructure has started to be fully recognized. Several powerful concepts have emerged to facilitate this vision, whether involving the human context whenever required or directly impacting user behavior and decisions. As these become the stepping stones to develop the IoT into a novel people-centric utility, this paper outlines a path to materialize this decisive transformation. We begin by reviewing the latest progress in human-aware wireless networking, then classify the attractive human–machine applications and summarize the enabling IoT radio technologies. We continue with a unique system-level performance characterization of a representative urban IoT scenario and quantify the benefits of keeping people in the loop on various levels. Our comprehensive numerical results confirm the significant gains that have been made available with tighter user involvement, and also corroborate the development of efficient incentivization mechanisms, thereby opening the door to future commoditization of the global people-centric IoT utility.**

I. Introduction and Vision

The Internet of Things (IoT) has undergone a fundamental transformation in recent decades: departing from the legacy radio frequency identification (RFID) technology of the 1980s and the wireless sensor networks (WSNs) of the 1990s, which were essentially siloed "connectivity islands" with limited interoperability, to its present form, which is becoming increasingly interconnected and heterogeneous. Today's IoT is already a fusion of numerous networked tools and appliances, equipped with advanced computational intelligence and rich communication capabilities. More broadly, the principles of contemporary IoT overlap with and permeate many adjacent domains, including mobile and pervasive computing, as well as robotics and cyber-physical systems – with applications ranging from smartphone-based social networking that reduces traffic and pollution in cities to mission-critical industrial automation that monitors and actuates over factory processes [1].

Beyond legacy embedded systems with constrained applicability, the emerging IoT solutions are becoming more open and integrated by adaptively combining sensors and actuators with actionable intelligence for automatic monitoring and control [2]. However, as a multitude of interconnected and intelligent machines communicate with each other and autonomously adapt to changing contexts without user involvement, the fact that present technology is made by humans and for humans is often overlooked [3]. Indeed, modern IoT systems are still widely unaware of the human context and instead consider people to be an external and unpredictable element in their control loop. Therefore, future IoT applications will need to intimately involve humans, so that people and machines could operate synergistically. To this end, human intentions, actions, psychological and physiological states, and even emotions could be detected, inferred through sensory data, and utilized as control feedback.

This concept, which is known as human-in-the-loop (HITL), opens the door to next-generation people-oriented IoT applications with more efficient and intuitive manipulation [4]. As users increasingly interact with such human-aware HITL systems, they may also become directly influenced by the control-loop decisions, thereby closing the loop. In fact, people may receive control input from the system in the form of suggestions and incentives (or even penalties) to diverge from their default behavior. Accordingly, human behavior may be impacted in either space (for example, the users are encouraged to move to a less congested location) or time (for example, the users are convinced to reduce their current data demand in case the network is overloaded); this is known as the "user-in-the-loop" (UIL) [5].

With UIL, often referred to as "layer 8", the space-time user traffic demand may be shaped opportunistically [6] and better matched with the actual resource supply from the people-centric wireless system. While HITL involves the user whenever human participation is desired or required and UIL extends the user's role beyond a traffic-generating and traffic-consuming black box [7], these trends must account for the fact that people are, in essence, walking sensor networks [8]. Indeed, a wide diversity of user-owned companion devices, such as mobile phones, wearables, connected vehicles, and drones may become an integral part of the IoT infrastructure. Hence, they can augment a broad range of applications, in which human context is useful, including traffic planning, environmental monitoring, mobile social recommendation, and public safety, among others. Therefore, we envision that – in contrast to past concepts where the user only assists the network to receive better individual service – future user equipment will truly merge with the IoT architecture to form a deep-fused human–machine system that efficiently utilizes the complementary nature of human and machine intelligence.

Should the IoT keep people in the loop, it has the potential to evolve into an integrated system-of-systems that may interact to form novel and unprecedented services [9]. For instance, the underlying people-centric sensor and actuator network may act as a utility – similar to electricity and water – creating important usable knowledge from vast amounts

V. Petrov, D. Moltchanov, S. Andreev, and Y. Koucheryavy are with the Department of Electronics and Communications Engineering, Tampere University of Technology, Finland.

K. Mikhaylov and M. Juntti are with the Centre for Wireless Communications, University of Oulu, Oulu, Finland.

H. Yanikomeroglu is with the Department of Systems and Computer Engineering, Carleton University, Canada.

G. Fodor is with Ericsson Research and Wireless@KTH, Sweden.

J. Torsner is with Ericsson Research, Finland.

This work was supported by the Academy of Finland and by the project TAKE-5: The 5th Evolution Take of Wireless Communication Networks, funded by Tekes. †V. Petrov is the contact author: vitaly.petrov@tut.fi

of data. Facilitated by this new global utility, different IoT devices that previously had nothing to do with each other may discover and start talking to one another, thereby augmenting the current talk-by-design approaches. This vision renders future IoT as a genuinely multi-application platform that can be materialized in the near future by means of candidate radio technologies [10], which we systematically review below. After describing a working people-centric IoT system as a case study, we also see how these technologies respond to the representative human user involvement models and quantify the resultant system-wide benefits.

II. Envisaged People-Centric IoT Applications

In light of the above, a contemporary perspective on the IoT expects it to soon become "the infrastructure of the information society" [11]. The very capable machines, such as sensors, actuators, robots, and vehicles, are already assisting humans in their daily lives. Explosive growth in the population of such connected objects leads to complex human–machine interactions that become increasingly frequent, facilitated by the HITL and UIL concepts. As these interactions intensify, many categories of people-centric IoT services emerge and are expected to be deployed over the following years:

- *Intention- and mission-aware services*. These services primarily reflect user's current intention or desire and assist by enabling, for example, situation-aware smart commuting for pedestrians, cyclists, and drivers of scooters, trucks, and other vehicles. This group of applications can help people in a variety of use cases, from highlighting the nearest available parking space on a vehicle's head-up display in urban areas to status reporting from stationary and wearable sensors on a display or augmented reality (AR) glasses in challenging environments, such as mines, construction sites, etc. In the near future, parking facilities may evolve into charging stations for electric vehicles that provide information not only about their availability, but also their characteristics, such as power, time to charge, and price.
- *Location- and context-aware services*. Another group of services is formed by location- and context-aware applications, such as those communicating alerts from environmental sensors (for example, "put on/take off your mask" when entering/leaving a polluted area). It may also include advice to go shopping when passing by a grocery store that has reasonably-priced goods from the user's shopping list. Many more of these alert-based services are envisioned to be deployed in the coming years, such as identifying slippery floors and low ceilings, notifying about forgotten trash and open taps when a user is about to leave the house, and many other examples.
- *Condition- and mood-aware services*. A deeper level of IoT penetration into people's lives can be achieved by integrating city/area infrastructure with personal medical and wellness devices. For instance, dietary restrictions could be applied on a menu when ordering food or a squad leader may be advised to give a break to a worker whose blood pressure has recently gone up. Other examples include smart vehicles that refuse to start if about to be operated by a drunk driver, gun safes not opening to an individual whose mental health is in question, and so on.

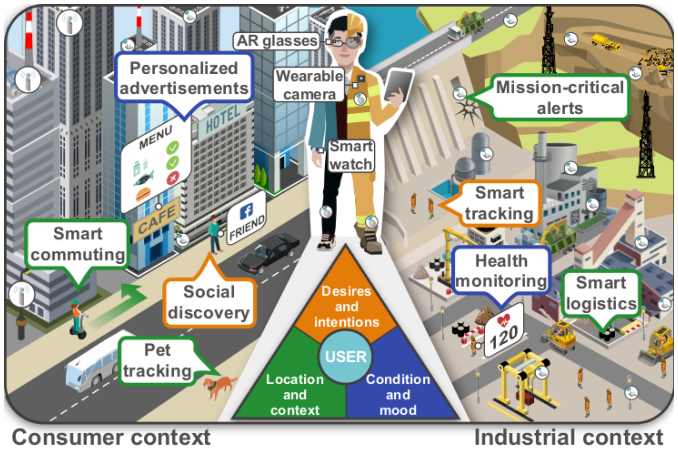

Fig. 1. Consumer and industrial contexts of people-centric IoT applications.

Summarizing the above examples of services, we note that depending on the environment the set of requirements and challenges to implement a particular application may vary considerably. For example, smart commuting in an urban area is very different from that inside a mine. To further offer a challenges-based grouping, we propose to differentiate between two major contexts: consumer and industrial (see Fig. 1). The former is characterized by the presence of numerous devices that are heterogeneous in terms of their communication means and ownership. Therefore, the major challenge in this context is to provide sufficient *scalability* of the deployed connectivity solution. On the contrary, the latter context is more challenging in terms of maintaining communication *reliability* due to more difficult propagation environments. At the same time, the system operator has more control over device population in such areas.

We continue by addressing how people-centric IoT applications are to be engineered, that is, which radio technologies need to be employed in particular scenarios and how they should be adjusted to ensure their usability in target operating conditions.

III. Review of Wireless Connectivity Options

Inspired by the above use cases, we review the contemporary IoT radio access technologies (RATs) and analyze them through the prism of their applicability for human–machine applications. Below we collect several of the emerging RATs that are well supported by the industrial community and thus have very high market potential to empower future HITL-/UIL-based applications.

1) Current technology diversity: The demand for developing RATs that focus specifically on the needs of machines was commonly understood in the mid to late 2000s. The low power consumption, affordable cost, high communication range, and capability to handle massive deployments of infrequently transmitting devices became the major requirements for these new solutions, which can be collectively referred to as Low-Power Wide-Area Networks (LPWANs). Development of such

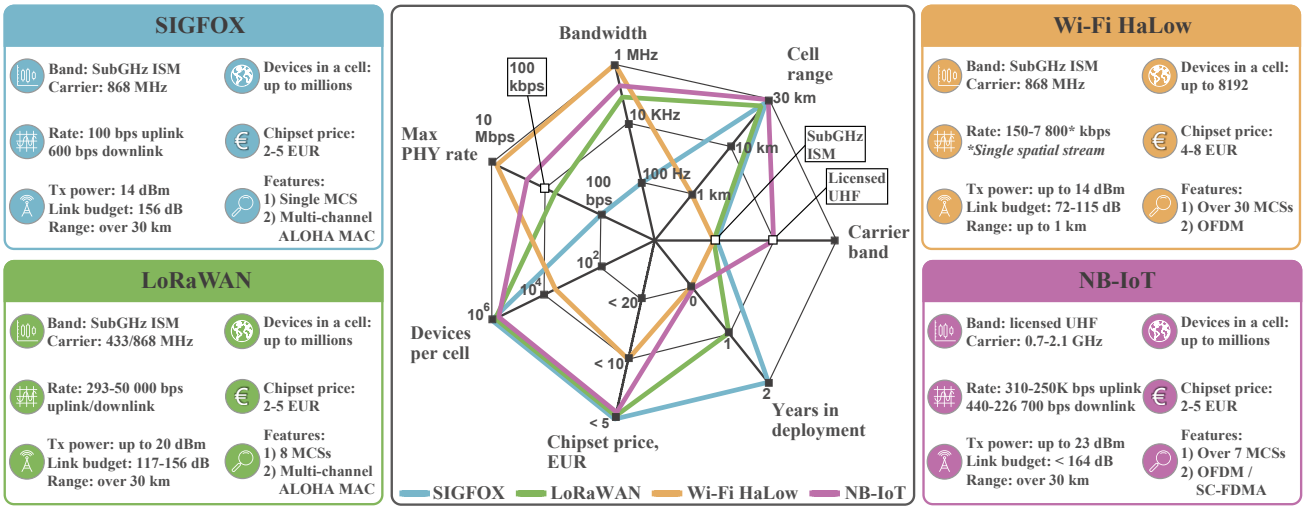

Fig. 2. Key characteristics of candidate radio technologies for advanced human–machine interaction.

technologies emerged and progressed in parallel within several leading standardization bodies, including IEEE 802.11 and 802.15 groups, ETSI, and 3GPP.

This resulted in the sheer diversity of today's LPWAN options, which comprise a number of standardized solutions (for example, extended coverage GSM, LTE-M, and NB-IoT, IEEE 802.15.4g and 802.15.4k, ETSI LTN, etc.), several (partially) open protocols pushed by alliances or special groups (for example, Weightless and LoRaWAN), as well as a few dozens of proprietary closed-specification technologies (for example, SIGFOX, RPMA, and WAVIoT). Across this diversity, the paths taken by the technology developers differ substantially. To offer illustrative examples, we consider two emerging LPWAN solutions, namely, SIGFOX and LoRaWAN. Both technologies operate in the sub-GHz license-exempt industrial, scientific, and medical (ISM) frequency bands and employ ALOHA-based channel access with frequency hopping, which places them under severe duty cycle restrictions. Topologically, both alternatives adhere to a cellular-like structure.

The SIGFOX solution operates with ultra-narrow band (UNB) signals: 100 Hz using Binary Phase-Shift Keying (BPSK) modulation and 600 Hz using Gaussian Frequency Shift Keying (GFSK) modulation for uplink and downlink, respectively. A short packet is sent on a random channel at a random time, thus implementing the so-called Random Frequency and Time Division Multiple Access (RFTDMA) scheme. The base stations listen over a sufficiently broad frequency band (dozens to hundreds of kHz) to detect the transmitted signals. Even though SIGFOX imposes harsh restrictions on the application payload of a radio frame (i.e., 12 bytes at most), the on-air time for each frame can exceed 1 second. Other substantial restrictions of this technology are the single modulation and coding scheme (MCS) and the limitation on the number of packets that can be sent to and from a device per day to 4 and 140, respectively.

In contrast to SIGFOX, LoRaWAN supports multiple MCSs. The mandatory MCSs are based on the Semtech's proprietary LoRa Spread Spectrum (SS) modulation derived from the Chirp SS modulation. The number of chirps to carry a single bit as well as the bandwidth of the channel that affects the duration of a single chirp can be adjusted to trade the on-air time for the transmission range. Additionally, the 50 kbps GFSK modulation may be enabled in some regions.

2) Machines talk Wi-Fi: The path followed by the IEEE 802.11 community is substantially different. The work on IEEE 802.11ah standard (also known as Wi-Fi HaLow) is rapidly approaching its final stage and early chipsets implementing this technology have already been announced. The solution targets the "high-end" devices with demanding performance requirements in terms of throughput and latency. To this end, 802.11ah delivers a non-backward compatible communication technology in the sub-1 GHz (S1G) ISM band, which is based on orthogonal frequency division multiplexing (OFDM) and features several dozens of different channel bandwidths as well as modulation and coding options.

One of the intrigues that remain today is how close the relationship between the HaLow and the conventional Wi-Fi systems will be. That is, whether HaLow continues on its own or be merged with the traditional Wi-Fi to form dual-mode devices similar to Bluetooth Smart: for example, dual technology enabled access points or Wi-Fi HaLow access points combined with a Wi-Fi client to be used in a smartphone.

3) 3GPP goes machine: Machine-type communication received considerable attention in 3GPP. Back in the early 2010s, 3GPP focused efforts on outlining the technology named LTE-MTC or LTE-M. Addressing the need for reduced cost and energy consumption while increasing the coverage range and the number of served devices per cell, the said technology has made so far several decisive steps. First, in LTE Release 12 the simplified Cat. 0 user equipment (UE) was introduced to operate with 20 MHz bandwidth and provide the peak data rates of 1 Mbps. Further, following the lines of the TR 36.888 document, a novel UE category named Cat. M1 (included into LTE Release 13) was designed to operate over the bandwidth of 1.4 MHz with reduced maximum transmit power, while still providing a substantial coverage extension.

Despite the immense progress in global LTE deployment, much area is still covered by the legacy GSM solutions, where cellular machine-type applications employ GPRS/EDGE for basic connectivity. Therefore, GSM is likely to continue playing its role in the IoT evolution due to its time-to-market and cost advantages. Recognizing this, 3GPP decided to "hone" GSM for the IoT market by providing a better alternative for the IoT traffic than GPRS. Named extended coverage GSM

(EC-GSM), it is implemented by defining new control and data channels mapped over legacy GSM as well as features sophisticated repetition and signal combining techniques. As a result, EC-GSM can theoretically provide cell capacity of up to 50,000 devices, while coverage may be extended by up to 20 dB at the sub-GHz carrier.

Finally, in September 2015, following the recent activities on the cellular IoT condensed in the TR 45.820 document, the work on the narrowband IoT (NB-IoT) has officially commenced. The new Cat. NB1 devices can be integrated into today's communication networks and enable UEs with about 10% complexity of that for Cat. 1. To achieve this goal, the bandwidth is reduced down to 180 kHz, OFDMA with 15 kHz subcarrier spacing is planned for the downlink, FDMA with 15 kHz subcarrier spacing and SC-FDMA with either 3.75 or 15 kHz between the subcarriers operate in the uplink. NB-IoT can offer three deployment options: standalone, in-band on the LTE carrier, or in the LTE guard band. The transport block sizes for both uplink and downlink have been reduced as well as many signaling procedures have been simplified. The standardization process behind NB-IoT was completed in June 2016 and the solution was included into LTE Release 13.

IV. Developed Evaluation Methodology

With the aim to characterize the tentative performance gains of various user involvement mechanisms, we concentrate on a representative urban use case that may correspond to several practical IoT applications (see Fig. 1): smart commuting, context-aware alerts, personalized advertisements, etc. The following text summarizes the considered deployment and network topology, the proposed user involvement strategies, and the details of our conducted simulation study.

1) Deployment parameters: In this work, we focus on a typical urban scenario over a 1 km^2 Manhattan grid deployment (see Fig. 3). A total of 100 city blocks are modeled. We consider two types of connected machines: stationary devices that represent smart road signs, environmental sensors, etc., and mobile wearable machines that offer healthcare, fitness, and well-being functionality. Following Google StreetView data on the density of road signs and pedestrians on a weekly basis in Manhattan, New York City, USA[1], the number of connected machines is set to its minimum feasible value of 20 pre-deployed devices per block and 30 wearable devices per block. This leads to 2000 stationary and 3000 moving machines across the entire simulation scenario. All of the modeled devices are deployed uniformly over the sidewalks.

We also model vehicles that participate in intense downtown traffic, where cars are driving along the streets with the constant speed of 30 km/h. The random inter-vehicle distance follows an exponential distribution with the mean of 3 m. At the intersections, we adopt the Manhattan mobility pattern: vehicles continue in the same direction with the probability of 0.5, while the chances to make a left or right turn are equal to 0.25. Pedestrians carrying wearable machines follow a similar mobility pattern with the speed of 5 km/h. To eliminate

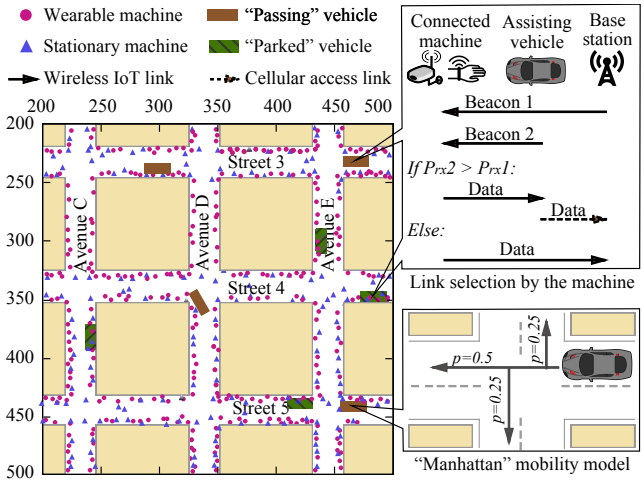

Fig. 3. Part of our representative urban IoT scenario: 300 m × 300 m out of 1 km × 1 km area is displayed.

border effects, a wraparound mechanism is implemented, such that any mobile device leaving the area of interest on a side immediately appears on the opposite side of the map.

2) Simulation details: The reported performance assessment has been performed with our custom-made system-level simulator, named WINTERsim[2], which has been extensively utilized recently for studying various IoT scenarios [12]. In the present study, we focus primarily on the uplink IoT traffic by modeling connectivity between a number of machines and the base station. To fairly compare the behavior of all the four considered RATs (i.e., to avoid overloading SIGFOX), the data transmissions are assumed to be regular and infrequent. More specifically, stationary machines communicate 10 B messages every 5 s, while moving devices send a 100 B update every minute, which corresponds to the typical sensing-based applications (e.g., assessing user's physical condition with a medical sensor).

From the connectivity perspective, machines communicate with their serving base station by default (termed *baseline* connectivity). Alternatively, if a better signal level is available, the devices may also connect to one of the femtocell relay stations deployed on the mobile vehicles across the tracking area (termed *assisted* connectivity). To characterize the effects of involving user-owned relaying cars for the "average" IoT device, "median-quality" connections to the base station are of particular interest. Therefore, in our scenario the base station is located at a certain distance from the center of the area to ensure the baseline average SINR level of 10 dB. As the properties of the considered LPWAN RATs vary significantly, four different simulation rounds have been completed, with the distances to the base station ranging from 520 m for Wi-Fi HaLow and up to 12 km for LoRaWAN.

Certain additional assumptions have been made to implement the LPWAN technologies of interest. Particularly, the hard message limitation for SIGFOX is relaxed. Then, LoRaWAN is restricted to operate exclusively over the 868 MHz band. The bandwidth of Wi-Fi HaLow, in its turn, is set to 1 MHz disregarding the 2 MHz and wider bandwidth options that offer higher data rates for the cost of reduced coverage.

[1]Particularly, Midtown East between 2nd and 3rd Avenue as well as 42nd and 57th Street is taken as an example

[2]WINTERsim system-level simulator: http://winter-group.net/download/

TABLE I
SIMULATION SETUP AND PARAMETERS

Parameter	Value
Message payload (stationary machine)	10 B
Message payload (moving machine)	100 B
Inter-arrival time (stationary machine)	5 s
Inter-arrival time (moving machine)	60 s
City block size	80 m
Street width	25 m
Scenario size (100 blocks)	1050 m × 1050 m
Total number of stationary machines	2000 / scenario
Total number of moving machines	3000 / scenario
Base station height	10 m
Car roof height	1.5 m
Deployment height of stationary machines	Uniform in $[0; 10)$ m
Deployment height of moving machines	1.5 m
Number of vehicles	1000 / scenario
Number of pedestrians	3000 / scenario
Speed of vehicles	30 km/h
Speed of pedestrians	5 km/h
Average inter-vehicle distance	3 m
Inter-vehicle distance distribution	Exponential
Mobility pattern of vehicles and pedestrians	"Manhattan"

Finally, only 7 key MCSs are supported for NB-IoT. The remaining radio technology related parameters are adopted from Fig. 2, while other important settings are summarized in Table I.

In our subsequent evaluation, we focus on the two key metrics: (i) SINR at the receiver and (ii) energy efficiency of machines, which is defined as the amount of data that has been reliably delivered from the IoT device to the base station related to the total energy consumed by the machine's radio interface.

3) Degrees of user involvement: As relaying vehicles in our IoT scenario are not owned by the operator but by the private users, the assisted connectivity option requires certain levels of user involvement. In this article, we study two different types of such involvement in the characteristic urban IoT deployment:

• *User involvement Type 1*. In this case, some of the driving vehicles share their communication capabilities and, if in close proximity, can opportunistically forward the IoT traffic from the machines to the application server. The deployed machines compare the connection quality to the base station with that to the nearest vehicle that is willing to assist, and transmit their following update over a better channel. In its turn, the assisting vehicle relays thus received IoT data to the application server over its on-board communication equipment.

• *User involvement Type 2*. As the user involvement Type 1 does not affect the vehicle mobility patterns, it can only offer opportunistic gains. On the contrary, the user involvement Type 2 suggests the car owner to temporarily park the vehicle close to the cluster of machines suffering from inadequate connection quality. For instance, as a reward the car owner may be offered a free or discounted parking permit in the said location (or free charging time in case of an electric vehicle). In this case, conforming vehicles are placed next to the centers of device clusters and start continuously forwarding the IoT data from the machines to the application server. Accordingly, a link with a better quality can be made available to the machines for longer time intervals. However, the user involvement Type 2 may request humans to deviate from their intended mobility patterns and the corresponding sources of motivation should thus be provided.

The achievable performance gains in terms of machine's energy efficiency for both types of user involvement are reported in subsection V-A, while a discussion on the nature and the origins of user involvement is offered in subsection V-B.

V. BENEFITS AND NATURE OF USER INVOLVEMENT

A. Achievable Performance Gains

To ease further exposition, we only account for the proportion of data transmitted by the machines. To this aim, we reasonably assume that the wireless connections from the privately-owned vehicles (which are neither battery- nor power-constrained) to the base station are reliable enough to guarantee the delivery of the relayed IoT data. As we address a massive machine-type deployment, we focus on the average performance gains by aggregating the effective measured values across individual IoT devices.

First, Fig. 4 reports on the average levels of SINR at the receiver (either the base station or the assisting vehicle, depending on the machine's current connectivity option). We learn that for all the IoT technologies in question the considered user involvement options notably outperform the baseline alternative. In particular, noticeable SINR gains are achieved already with a few assisting vehicles. The SINR gains for the SIGFOX solution are the highest, while the other three LPWAN protocols behave similarly as the degree of human assistance increases. The explanation behind these results is in the fact that SIGFOX operates in a full-power regime with a single data rate, while the rest RATs allow the machines to dynamically select their transmit power and MCS. The results for the user involvement Type 2 are, on average, several dB higher than the corresponding numbers for the user involvement Type 1, which confirms that careful positioning of assisting nodes is preferred to their opportunistic placement.

Further, Fig. 5 illustrates the impact of the studied user involvement types on the average energy efficiency of the communicating machines. This parameter offers insights into how far the battery life could be extended and, consequently, how much the operating costs could be reduced. The reported energy efficiency may be converted into battery lifetime subject to appropriate scaling that depends on the type of machine (for example, a small temperature sensor or a self-driving goods delivery robot) and its energy consumption profile. In the figure, we first observe a marginal impact of user involvement on the energy efficiency of SIGFOX. Albeit substantial improvements in the SINR are confirmed by Fig. 4, the single transmit power level in SIGFOX limits the potential benefits of the assisting vehicles. The observed gain of only several percent is due to a slightly increased level of reliability.

On the contrary, the other three RATs demonstrate considerable energy efficiency improvements (for example, up to 4 times greater energy efficiency for NB-IoT with user involvement Type 2). The observed growth is not only due to better channel conditions (which are similar for the said technologies in the same scenario), but also due to the flexibility of the IoT RATs themselves. Specifically, NB-IoT outperforms its counterparts owing to more freedom in the MCS and transmit

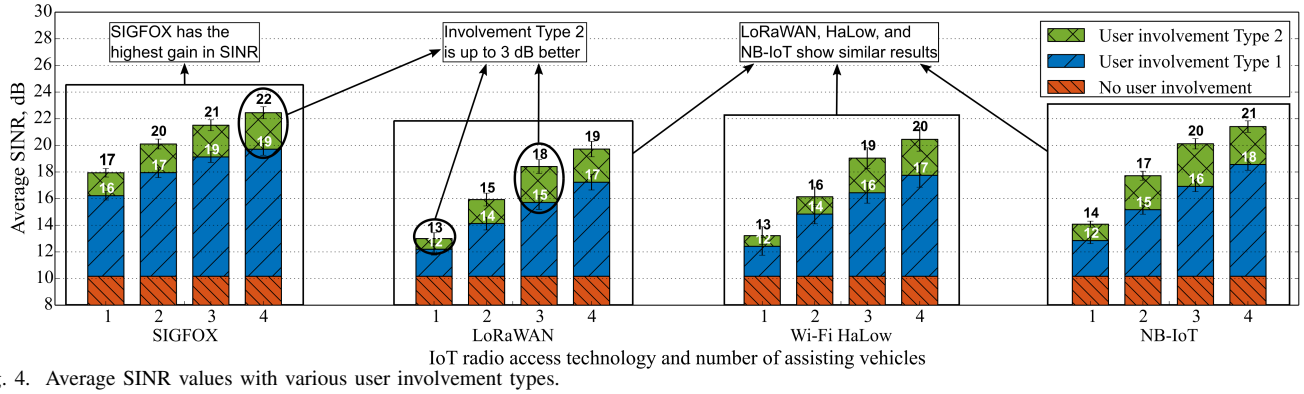
SIGFOX has the highest gain in SINR
Involvement Type 2 is up to 3 dB better
LoRaWAN, HaLow, and NB-IoT show similar results
User involvement Type 2
User involvement Type 1
No user involvement
Average SINR, dB
SIGFOX
LoRaWAN
Wi-Fi HaLow
NB-IoT
IoT radio access technology and number of assisting vehicles

Fig. 4. Average SINR values with various user involvement types.

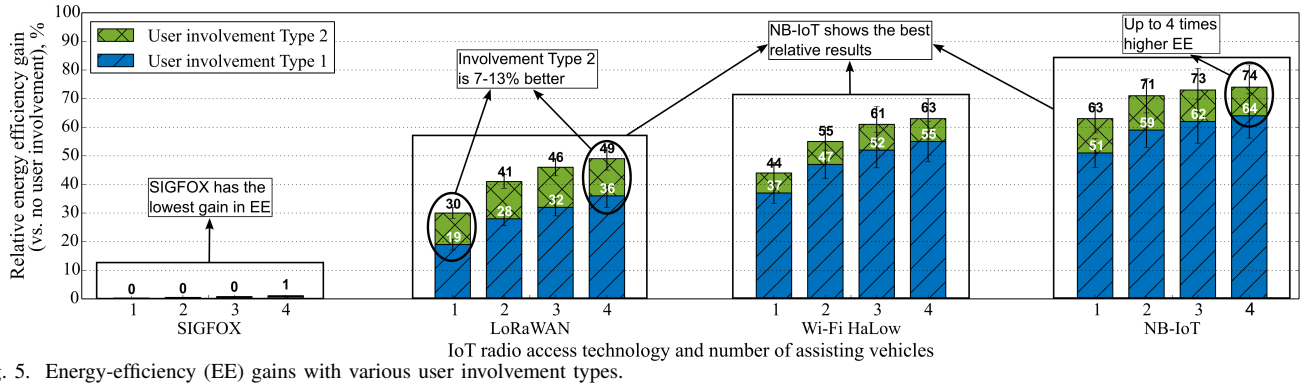
User involvement Type 2
User involvement Type 1
SIGFOX has the lowest gain in EE
Involvement Type 2 is 7-13% better
NB-IoT shows the best relative results
Up to 4 times higher EE
Relative energy efficiency gain (vs. no user involvement), %
SIGFOX
LoRaWAN
Wi-Fi HaLow
NB-IoT
IoT radio access technology and number of assisting vehicles

Fig. 5. Energy-efficiency (EE) gains with various user involvement types.

power selection. In its turn, Wi-Fi HaLow shows benefits over LoRaWAN in the investigated scenario. This is because the performance of Wi-Fi HaLow (as a shorter range radio compared to LoRaWAN: 1 km vs. 30 km) is much stronger affected by a decrease in the effective link length from hundreds to tens of meters. Another reason for moderate performance improvement of LoRaWAN is the presence of non-adaptive contention-based channel access mechanism, which in massive IoT scenarios leads to collisions even at low traffic volumes.

In summary, the considered ways of user involvement may lead to notable improvements in the IoT service reliability (due to higher SINR) as well as energy efficiency of networked machines. Hence, people in the loop assist the operator in reducing both the capital investments and the operating costs. In the remainder of this text, we discuss several approaches to materializing this gain in practice.

B. User Incentivization Options

A crucial component of the considered system operation is related to user's willingness to share own resources while helping to improve the IoT network performance. Even though the concept of user involvement for resource provisioning has been discussed for years in the context of various mobile systems (examples include client relaying, device-to-device communications [13], etc.), these mechanisms have not been implemented widely as of yet. This is in part owing to conservative human nature that resists new models of service provisioning, especially when the process requires involving own resources without immediately perceived benefits. We envision that appropriate user incentivization mechanisms need to be natively integrated into the envisioned IoT infrastructure. In systems with regular topologies and central control – that emerge primarily in the industrial context – there always remains an opportunity to enable resource sharing by default. On the contrary, strictly enforcing resource sharing in consumer scenarios may lead to customer dissatisfaction and therefore clever user-friendly incentivization mechanisms are required.

Developing appropriate incentivization schemes for the consumer context, one has to mediate between dissimilar interests of at least five major stakeholders: (i) owners of the predeployed machines; (ii) owners of the wearable machines; (iii) owners of the assisting vehicles; (iv) car manufacturers, and (v) system operators. While the former two groups secure immediate benefits from the considered operation (see the previous subsection), the remaining three players need to first share their resources. The motivation for a car manufacturer can be straightforward as the support of the discussed mechanism increases the amounts of data collected and processed by the vehicles, thus making them even more intelligent. This also improves the market value of the produced vehicle to the customers who are willing to collect certain rewards in the future. In its turn, the potential direct or indirect compensation can become a sufficient motivation for the car owners, especially taking into account that the actual expenses are minimal – different from the smartphone, a relative energy consumption by the assisting vehicle is negligible. Finally, the network operators can expect a distinct benefit due to higher satisfaction levels by the first two groups of their customers, which is directly converted into more connected machines, mobile subscriptions, etc.

However, to deploy the discussed mechanisms in practice, the operators need to carefully "moderate the dialog" between all of the parties, where the most non-trivial aspect is to engage the owners of vehicles. In order to make it happen, a service provider has to offer (i) a type of compensation for the car owners and (ii) a corresponding billing mechanism. For the former, there is a set of options including those directly related to the operator–user monetary relationships, such as decreased monthly rates and/or extended subscription periods. Considering the operator's network as a collection

of resources, one could further identify a set of additional factors, including a priority in service, guaranteed rate during congestion, etc. A universal mapping scheme that translates the collected bonus points into the user benefits must also be provided. It is crucial that such a mapping is made to operate on-line as users may not remember at the end of the month what they are rewarded for exactly.

Further, a billing mechanism is a necessary component of the incentive-driven services. The set of requirements imposed on its choice includes robustness to a loss of control connection, high levels of security during authorization and encryption of all the exchanged billing-specific information, operational accuracy, as well as simplicity. There are two fundamental paths to implementing an efficient billing algorithm. In a fully centralized system, the operator becomes responsible for all the phases, such as collection, storage, and interpretation of information on the contributions by the users. To alleviate the cases of false notification by malicious participants, each session has to be authorized and information needs to be collected from both endpoints. The main benefit of centralized billing is in full control over the billing-related data. However, this approach may see difficulty in handling the situations when network infrastructure becomes temporarily unavailable. On top of this, such a scheme is vulnerable to specific malicious attacks, e.g., two participants exchanging useless data with the aim of earning more benefits.

An alternative to the above is a network-assisted decentralized billing solution [14]. There are numerous practical approaches, where decentralized incentivization mechanisms have been successfully deployed. A famous example is peer-to-peer file sharing and streaming services. There, reciprocity-based mechanisms, where the peers maintain and share the information about other peers' contributions, remain a popular method. In wireless systems, where only two entities are involved into the process of information exchange at a given time, one could extend the reciprocity-based algorithms to design network-wide reputation metrics. In such a scheme, network assistance is still required for authentication purposes, but the actual billing process could be made decentralized. However, fully decentralized mechanisms may not be feasible in commercial consumer systems due to a lack of stability.

VI. Conclusions

The unprecedented proliferation of the interconnected autonomous machines increases drastically the intensity and the depth of human–machine interactions. On the one hand, this introduces a wide range of novel challenges with respect to how the individual procedures, technologies, as well as the overall working services and applications should be engineered – mindful of unique capabilities and limitations of humans and machines alike. On the other hand, this brings along an excellent opportunity to jointly engage both sides to take advantage of rich mutual collaboration and shield any weaknesses with each other's strengths. In this article, we have confirmed that this is valid for people as well as connected machines.

Our novel analysis of an illustrative consumer IoT scenario for different user involvement levels and across four perspective radio technologies demonstrates that even moderate human assistance makes a decisive difference and unlocks significant energy savings for networked machines. Importantly, our offered numerical results show that while each of the addressed wireless solutions does benefit from keeping people in the loop, the actual gains for every system profile depend on the flexibility of the underlying access protocol. At the same time, the costs associated with engaging the human context may include certain incentivization-related expenditures, which calls for careful planning of user involvement in future people-centric IoT infrastructure. Ultimately, the rapidly maturing human-aware IoT ecosystem may become a new global utility, thus "disappearing" in the "fabric of everyday life" [15], while enabling a plethora of next-generation applications and services.

Acknowledgment

The authors are particularly grateful to Yngve Selén (Wireless Access Networks Department, Ericsson Research) for the insightful comments that allowed to improve this paper.

References

[1] D. Nunes, J. S. Silva, C. Herrera, and F. Boavida, *Human-in-the-Loop connectivity management in smartphones*, pp. 159–170. Springer International Publishing, 2016.

[2] C. Evers, R. Kniewel, K. Geihs, and L. Schmidt, "The user in the loop: Enabling user participation for self-adaptive applications," *Future Generation Computer Systems*, vol. 34, pp. 110–123, May 2014.

[3] F. Boavida *et al.*, *People-Centric Internet of Things—Challenges, Approach, and Enabling Technologies*. Springer International Publishing, 2016.

[4] D. S. Nunes, P. Zhang, and J. S. Silva, "A survey on Human-in-the-Loop applications towards an Internet of All," *IEEE Communication Surveys & Tutorials*, vol. 17, pp. 944–965, Second quarter 2015.

[5] R. Schoenen, H. Yanikomeroglu, and B. Walke, "User in the Loop: Mobility aware users substantially boost spectral efficiency of cellular OFDMA systems," *IEEE Communications Letters*, vol. 15, pp. 488–490, May 2011.

[6] R. Schoenen and H. Yanikomeroglu, "User-in-the-Loop: Spatial and temporal demand shaping for sustainable wireless networks," *IEEE Communications Magazine*, vol. 52, pp. 196–203, February 2014.

[7] M. Mirahsan, R. Schoenen, H. Yanikomeroglu, G. Senarath, and N. Dung-Dao, "User-in-the-Loop for HetHetNets with backhaul capacity constraints," *IEEE Wireless Communications*, vol. 22, pp. 50–57, October 2015.

[8] B. Guo, Z. Wang, Z. Yu, Y. Wang, N. Y. Yen, R. Huang, and X. Zhou, "Mobile crowd sensing and computing: The review of an emerging human-powered sensing paradigm," *ACM Computing Surveys*, vol. 48, pp. 1–31, Aug. 2015.

[9] J. A. Stankovic, "Research directions for the Internet of Things," *IEEE Internet of Things Journal*, vol. 1, pp. 3–9, Feb 2014.

[10] M. R. Palattella, M. Dohler, A. Grieco, G. Rizzo, J. Torsner, T. Engel, and L. Ladid, "Internet of Things in the 5G era: Enablers, architecture and business models," *IEEE Journal on Selected Areas in Communications*, vol. 34, no. 3, pp. 510–527, 2016.

[11] ITU, "Internet of things global standards initiative," tech. rep., July 2015.

[12] S. Andreev, O. Galinina, A. Pyattaev, M. Gerasimenko, T. Tirronen, J. Torsner, J. Sachs, M. Dohler, and Y. Koucheryavy, "Understanding the IoT connectivity landscape: A contemporary M2M radio technology roadmap," *IEEE Communications Magazine*, vol. 53, pp. 32–40, September 2015.

[13] G. Fodor, E. Dahlman, G. Mildh, S. Parkvall, N. Reider, G. Mikls, and Z. Turnyi, "Design aspects of network assisted device-to-device communications," *IEEE Communications Magazine*, vol. 50, pp. 170–177, March 2012.

[14] R. Schoenen, H. U. Sokun, and H. Yanikomeroglu, "Green cellular demand control with User-in-the-Loop enabled by smart data pricing using an effective quantum (eBit) tariff," in *Proc. of the IEEE Vehicular Technology Conference (IEEE VTC-Fall), Invited paper*, Sep. 2016.

[15] M. S. Corson, R. Laroia, J. Li, V. Park, T. Richardson, and G. Tsirtsis, "Toward proximity-aware internetworking," *IEEE Wireless Communications*, vol. 17, pp. 26–33, December 2010.